

AWS Data Transfer Services

Deep Dive on Large-scale Migration and
Edge Compute with the AWS Snow Family

Ian Perez Ponce, Sr. Business Development Manager

What's driving storage relevance?



Artificial
Intelligence



Natural Language
Processing



Internet of
Things



Information
Assets



Data Trust
Frameworks

Unlimited Storage Scale

AWS cloud storage is core

Gives you unique scale...

- ✓ Greatest reliability
- ✓ Broad security and compliance
- ✓ Diverse portfolio
- ✓ Fastest innovation



Data matters at
any scale

Building cloud-native applications or **migrating** existing ones to AWS...

- ✓ Advanced developer tools
- ✓ Experienced consulting and support
- ✓ Methodical migration services
- ✓ **The most data movement services**

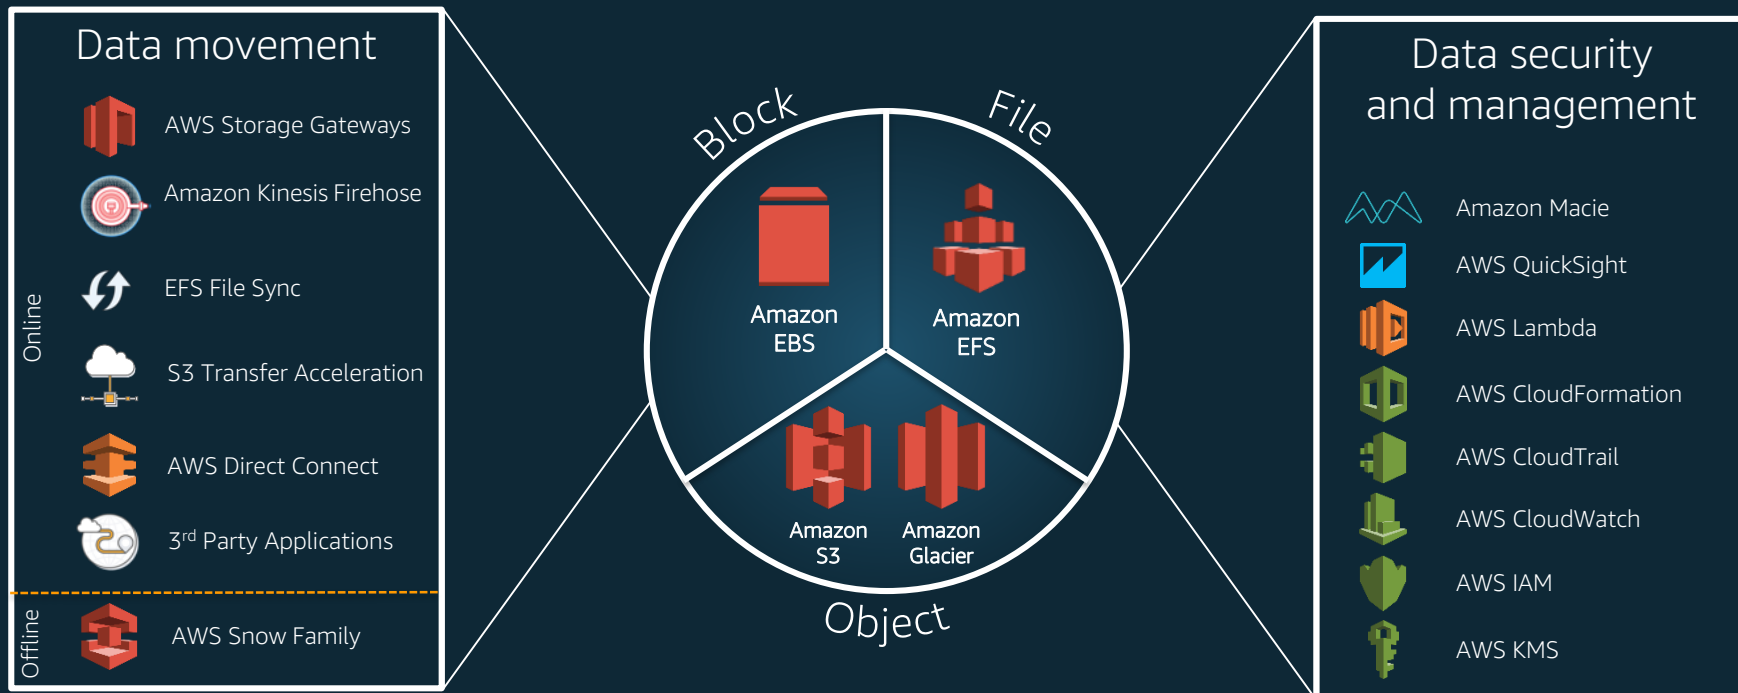
Yielding bigger insights...

- ✓ Most big data & data lake deployments
- ✓ Most managed databases
- ✓ Easiest data warehousing
- ✓ Singular query-in-place analytics

Helping you innovate and transform faster...

- ✓ Artificial Intelligence
- ✓ Deep Learning / Machine Learning
- ✓ IoT

The broadest range of storage services



More ways to move data to/from/within the cloud

Networks

AWS Direct Connect



A private connection between your data center, office, or colocation environment and AWS.

Amazon S3 Transfer Acceleration



Up to 300% faster transfers into and out of Amazon S3. Ideal when working with long geographic distances.

Amazon EFS File Sync



Up to 5x faster file transfers than open-source tools. Ideal for migrating data into Amazon EFS or moving between cloud file systems.

Amazon Kinesis Data Firehose



Capture, transform, & load streaming data into Amazon S3 for use with Amazon business intelligence and analytics tools.

Physical

AWS Snow family

(Snowball, Snowball Edge, Snowmobile)



Secure, physical transport appliances that move up to exabytes of data into and out of AWS.

Hybrid

APN competency partners



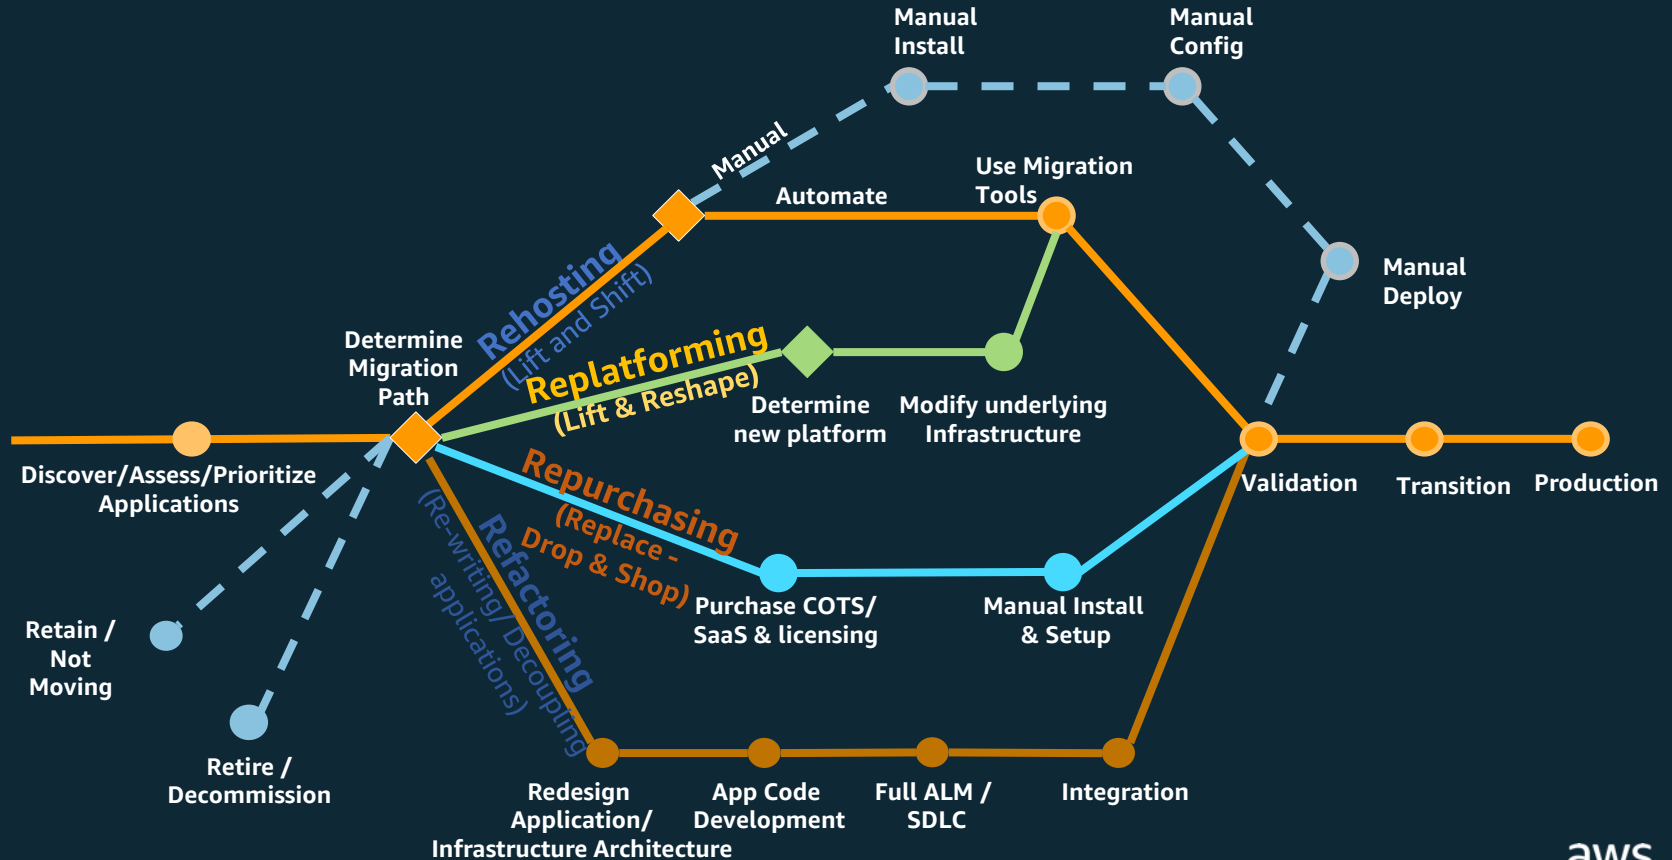
Integrations between third-party vendors and AWS services. Ideal for leveraging existing software licenses and skills.

AWS Storage Gateway




Hybrid storage that seamlessly connects on-premises applications to AWS storage. Ideal for backup, DR, bursting, tiering, or migration.

Application migration patterns



Don't overlook data migration fundamentals



Resources to help you migrate

AWS MIGRATION HUB



A single location to track the progress of application migrations across AWS and partner solutions

AWS MIGRATION SERVICES

- AWS Application Discovery Service
- AWS Server Migration Service
- AWS Database Migration Service
- AWS Schema Conversion Tool
- VMware Cloud on AWS

AWS DATA TRANSFER

- S3 Transfer Acceleration
- AWS Storage & File Gateway
- AWS Snowball & Snowmobile
- AWS Direct Connect
- Amazon Kinesis Firehose

Inventory



Business Case



Deep Discovery & Planning



App Dependency Mapping



Workload & Data Migration



Validation



Partner migration tools are vetted by AWS

Found on aws marketplace

Large-scale offline data transfer & edge processing with the AWS Snow family

What is Snowball and why did we build it?

Moving large volumes of data over the internet can take years – we ship secure physical devices to you to transfer your data at the source before shipping it back for bulk import to the cloud.

The cloud is not always accessible from remote locations where connectivity is limited or intermittent – deploy ruggedized devices at the edge with local storage and compute capacity to process data without network dependencies.

Traditional shipping of conventional hard drives is laborious and error prone – our E-Ink shipping label and chain of custody tracking simplifies logistics at scale.



AWS Snowball



AWS Snowball Edge



AWS Snow Family



AWS Snowball

- 50 or 80TB storage capacity
- 10GE networking
- Data encryption end-to-end
- Rugged 8.5 G impact case
- Rain and dust resistant



AWS Snowball Edge

- Similar rugged design to Snowball with encryption
- 100TB storage capacity
- 10/25/40GE networking
- AWS Greengrass for local compute, messaging and caching
- EC2/AMI support for edge compute *New*



AWS Snowmobile

- Exabyte-scale storage in a 45ft container
- Data encryption end-to-end
- Dedicated security personnel
- GPS tracking, alarm monitoring, 24/7 surveillance, and optional additional security

Specifications and pricing details

Snowball

Snowball Edge

Snowmobile

Storage	50 or 80 TB	100 TB	100 PB
Interfaces	CLI, S3 SDK	S3, NFS ^{v4}	NFS ^{v4}
Network	10G	10/25/40G	40/100G
Power	200W	400W	~350kW
Compute	NA	m4.4xl (equivalent)	NA
Job Fee (Import)	\$200 or 250 per job	\$300 per job	\$0.005/GB per month
Daily Rate	\$15/day after 10 th day	\$30/day after 10 th day	NA

Commercially available in 16 regions.

Do you need the truck?




< 10PB



> 10PB

Snowball import workflow



Create a job

In transit to you



Connect the Snowball

Delivered to you



Copy data to the Snowball

Delivered to AWS



Your data is loaded into your bucket!

At AWS

Job created

(~10 days end-to-end)

Job completed

Snowball key features

S3 compatible
SDK adapter



Snowball CLI
tools



Automated small
file batching



Embedded E-Ink
display for shipping



Chain of custody
tracking



SOC, PCI, & HIPAA
compliant



Storage density
alternatives

Snowball Edge unique features

S3 Compatible
Endpoint



NFS File
Interface




Line rate data
transfer



Storage and
compute cluster



AWS Greengrass and
Lambda support



EC2 Compute
Instances support

New



Hardware enabled
data encryption

Who's using Snow* today?

Customers

ISV Partners



BuzzFeed

Capital One

DigitalGlobe

drive.ai

EOSS

Etsy

Fender

HESS

healthgrades

mercado Libre.com

MERCK

NBC UNIVERSAL

NETFLIX

Pfizer

photobox

shutterstock

Symantec

trulia

WAZEE
BIG-TAIL

COMVAULT

CloudBerry Lab

druva

NetApp

wanDISCO

And for what use cases?

Wholesale data
center migration

Backup seeding

Assisted database migration

NAS appliance
archival

Data lake creation

Active archive

Edge computing – when migrating your data over distance isn't an option.

Where is edge computing in effect?

Various industries ranging from research to transportation necessitate edge computing capabilities to help cope with environments where **data generation is decentralized**, **data volumes are significant** and **network connectivity is either inaccessible or intermittent**.



Research



Industrial



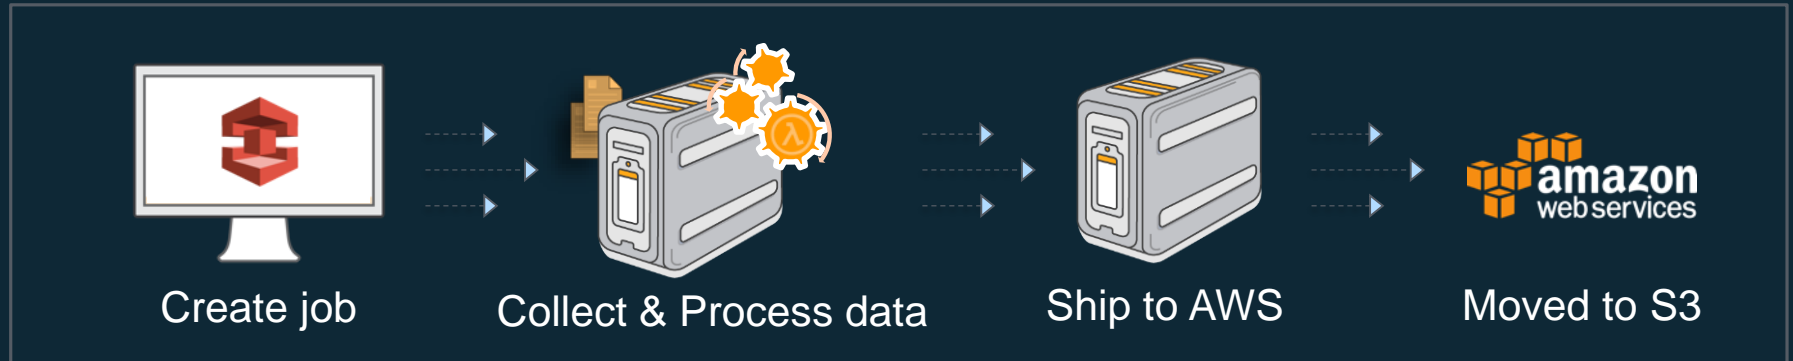
Healthcare



Transportation

Going beyond bulk data migration

- Snowball Edge supports the ability to capture, store and process data from variable sources deployed at the edge or in traditional data center environments
- Embedded compute running Greengrass or EC2 compute instances enables running a mix of serverless and traditional Linux or Windows based application workloads to satisfy various use cases for processing data at the edge



EC2 Compute Instances on Snowball Edge

New!

Service features announced

- Support for **EC2 sbe1 instances**
- Three AMIs available on AWS Marketplace
- SBE1 instances feature **1.8 GHz Intel Xeon D processors**, up to 24 vCPUs, 32 GiB of memory
- Up to 1TB disk volumes
- **1 and 3-year discounted pricing options**

Supported EC2 instance sizes

Model	vCPU	Mem (GiB)
sbe1.small	1	1
sbe1.medium	1	2
sbe1.large	2	4
sbe1.xlarge	4	8
sbe1.2xlarge	8	16
sbe1.4xlarge	16	32




Snowball
Edge

Snowball Edge clustering for scale

Clustering Features

- Scale from 5 to 10 nodes in a Snowball Edge cluster
- Clusters support aggregate storage and compute with leaderless nodes
- Increase durability for S3 object storage on-premises
- Easily swap nodes for maintenance



AWS
Snowman Cluster



EC2 Instances on Snowball Edge pricing

[AWS Snowball Edge](#)

[Overview](#)

[Features](#)

[Pricing](#)

[Getting Started](#)

[Resources](#)

[FAQs](#)

Pricing for Data Transfer Plus EC2 Compute Instances (varies by region)

Region	On-Demand Job Fee	On-Demand Extra Day Fee	1 year Commitment Upfront Charge*	3 years Commitment Upfront Charge*
US East (Northern Virginia, Ohio)	\$500	\$50	\$15,330 (\$42 per day)	\$38,325 (\$35 per day)
US West (Northern California, Oregon)	\$500	\$50	\$15,330 (\$42 per day)	\$38,325 (\$35 per day)
AWS GovCloud (US)	\$500	\$50	\$15,330 (\$42 per day)	\$38,325 (\$35 per day)
Canada (Montreal)	\$800	\$80	\$18,980 (\$52 per day)	\$49,275 (\$45 per day)
South America (São Paulo)	\$1200	\$120	\$36,500 (\$100 per day)	\$87,600 (\$80 per day)
EU (Ireland, Frankfurt)	\$800	\$80	\$21,900 (\$60 per day)	\$49,275 (\$45 per day)
EU (London)	\$800	\$80	\$21,900 (\$60 per day)	\$49,275 (\$45 per day)
EU (Paris)	\$800	\$80	\$21,900 (\$60 per day)	\$54,750 (\$50 per day)
Asia Pacific (Tokyo)	\$800	\$80	\$21,900 (\$60 per day)	\$54,750 (\$50 per day)
Asia Pacific (Singapore)	\$800	\$80	\$21,900 (\$60 per day)	\$54,750 (\$50 per day)
Asia Pacific (Sydney)	\$800	\$80	\$23,725 (\$65 per day)	\$60,225 (\$55 per day)

* Note: Upfront fee for full 1 year and 3 years results in per day discount. Please contact the AWS Sales team to obtain this pricing.

Problem

HealthSuite, our cloud-enabled connected health ecosystem of devices, applications and digital tools supported by AWS, is required to support over 14,000 global deployments, regardless of network access constraints.

Solution

Snowball Edge with EC2 and Lambda provides a viable platform for edge compute where common AWS services can be used to support medical imaging and optical scanning machines in use across geo-dispersed healthcare facilities.

Impact

Snowball Edge enables Philips Healthcare to extend the innovative capabilities of HealthSuite managed devices, even when there is limited or network support.



GE Oil & Gas



Problem

GE Oil & Gas uses pipeline inspection gauges (“PIGs”) to survey complex pipeline infrastructure for routine maintenance and operations which results in the need to collect and process data in often remote areas.

Solution

Initially relying on Snowball to capture PIG instrument data for import and analysis, Snowball Edge with EC2 now becomes a viable platform for field-level data collection and processing without having to incur a roundtrip to the cloud for analysis.

Impact

Field operations can now generate meaningful insights closer to the source where maintenance actions can be performed to increase operational efficiency.



Problem

With a growing fleet of research vessels for ocean ecosystem research, OSU requires constant collection of raw data for processing, analysis and archival.

Solution

With AWS Snowball Edge, OSU can now collect 100 TB of data with no intermediate steps and analyze imagery immediately using the onboard compute capabilities. This enables deeper analysis with the option to upload all raw data to the cloud by simply shipping the AWS Snowball Edge device back.

Impact

What used to take weeks to months for data ingest and processing using conventional disk drive technology and \$4M+ in local infrastructure investment is now a thing of the past.

Resources to get started

Snowball

Learn more at:

<https://aws.amazon.com/snowball/>

Webinars:

<https://aws.amazon.com/snowball/resources/>

10-minute tutorials:

https://aws.amazon.com/snowball/getting-started/#10-minute_tutorials

Getting started:

<https://console.aws.amazon.com/importexport/home>

Snowball Edge

Learn more at:

<https://aws.amazon.com/snowball-edge>

Videos:

<https://aws.amazon.com/snowball-edge/resources/>

Getting started:

<https://console.aws.amazon.com/importexport/home>

Thank you!